Freeform Search

Database: US Pre-Grant Publication Full-Text Database; US Patents Full-Text Database; US OCR Full-Text Database; EPO Abstracts Database; JPO Abstracts Database; Derwent World Patents Index; IBM Technical Disclosure Bulletins

Term: (first adj2 bus) same l69 same (partition or portion or region or section or bank) same (memory or cache)

Display: 10 **Documents in Display Format:** - **Starting with Number** 1

Generate: ○ **Hit List** ◉ **Hit Count** ○ **Side by Side** ○ **Image**

Search Clear Interrupt

Search History

DATE: Sunday, April 08, 2007 **Purge Queries** Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
	DB=PGPB,USPT,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ		
L70	(first adj2 bus) same l69 same (partition or portion or region or section or bank) same (memory or cache)	1	L70
L69	(exclusive or dedicated or private) adj3 (access or accessing)	11135	L69
L68	L67 and l66	4	L68
L67	(exclusive or dedicated or private) same (access or accessing)	82346	L67
L66	(common or shared) adj2 l37	30	L66
L65	L64 and l38 and l63	71	L65
L64	l37 or l41	2639	L64
L63	(dedicated or exclusive or private or (non adj2(shared or common))) adj4 (region or portion or section or partition or bank or access or accessing)	25831	L63
L62	l41 and l56	50	L62
L61	L59 same l58 and l37	0	L61
L60	L59 same l58	54	L60
L59	l25 adj5 processor adj5 (memory or cache)	135	L59
L58	l24 adj5 processor adj5 (memory or cache)	234	L58

L57	L56 same l37 same (processor or agent)	24	L57
L56	(parallel or simultaneous or concurrent) adj5 (access or accessing) adj5 (memory or cache)	3227	L56
L55	l37 same l26 same l27	13	L55
L54	l48 same (workload) same (processor or agent or cpu)	65	L54
L53	l52 and l51	15	L53
L52	l37 and l46	121	L52
L51	l37 and l45	64	L51
L50	L48 same l37	43	L50
L49	L48 and l37	296	L49
L48	(partition or partitioning or allocate or allocation or allocating or divide or dividing) adj4 (memory or cache)	33468	L48
L47	l37 same l32	124	L47
L46	(shared or common) adj4 (region or portion or section or partition or bank)	67718	L46
L45	(dedicated or exclusive or private or (non adj2(shared or common))) adj4 (region or portion or section or partition or bank)	13199	L45
L44	l37 and l12 and l13	2	L44
L43	l41 and l12 and l13	1	L43
L42	L41 and l39	0	L42
L41	711/149.ccls.	393	L41
L40	L39 and l37	0	L40
L39	l32 same l12 same l13	71	L39
L38	((multi adj2 processor) or (multiple adj2 (processor or agent)))	46597	L38
L37	((multi adj2 port) or (multiple adj2 port)) adj5 (memory or cache)	2339	L37
L36	L35 and l34	5	L36
L35	l25 same l13	8	L35
L34	l24 same l13	9	L34
L33	L32 and l30	8	L33
L32	(partition or partitioning or allocate or allocation or allocating or divide or dividing) same (memory or cache)	105982	L32
L31	partition	361687	L31
L30	L27 and l26 and l25 and l24 and l29	12	L30
L29	l13 same (memory or cache)	2244	L29
L28	L27 and l26 and l25 and l24 and l13	25	L28
L27	second adj2 (processor or agent)	40103	L27
L26	first adj2 (processor or agent)	38258	L26
L25	second adj2 bus	19990	L25
L24	first adj2 bus	19804	L24
L23	6965974.pn.	2	L23
L22	L21 and l20	4	L22
L21	bus same processor same l18	4	L21
L20	bus same processor same l17	4	L20

L19	(first adj2 bus) same processor same l17	0	L19
L18	second adj2 l13	164	L18
L17	first adj2 l13	173	L17
L16	l14 same l13 same l12	12	L16
L15	l14 and l13 and l12	233	L15
L14	(dynamically or (on adj2 fly)) same (changing or altering or alter or change or partition or partitioning or allocate or allocation or allocating or divide or dividing)	61847	L14
L13	(dedicated or exclusive or private or (non adj2(shared or common))) adj4 (region or portion or section or partition)	12227	L13
L12	(shared or common) adj4 (region or portion or section or partition)	65094	L12
L11	L10 and l8	4	L11
L10	l3 same processor	51	L10
L9	L8 and l3	7	L9
L8	L7 and l6	500	L8
L7	(partition or paritioned or partitioning).ti,ab.	224228	L7
L6	cache.ti,ab.	40875	L6
L5	L4 and l3	9	L5
L4	(dynamically or (on adj2 fly)) same (changing or altering or alter or change) same (partition or partitioning or allocate or allocation or region)	4258	L4
L3	l1 same l2 same (memory or cache)	153	L3
L2	(shared) adj5 (partition or region or partitioning or partitioned or divide or divided or allocate or allocating or allocation)	7906	L2
L1	(exclusive or dedicated or (non adj shared) or private) adj5 (partition or region or partitioning or partitioned or divide or divided or allocate or allocating or allocation)	6653	L1

END OF SEARCH HISTORY

Ruby Tuesday Restaurant

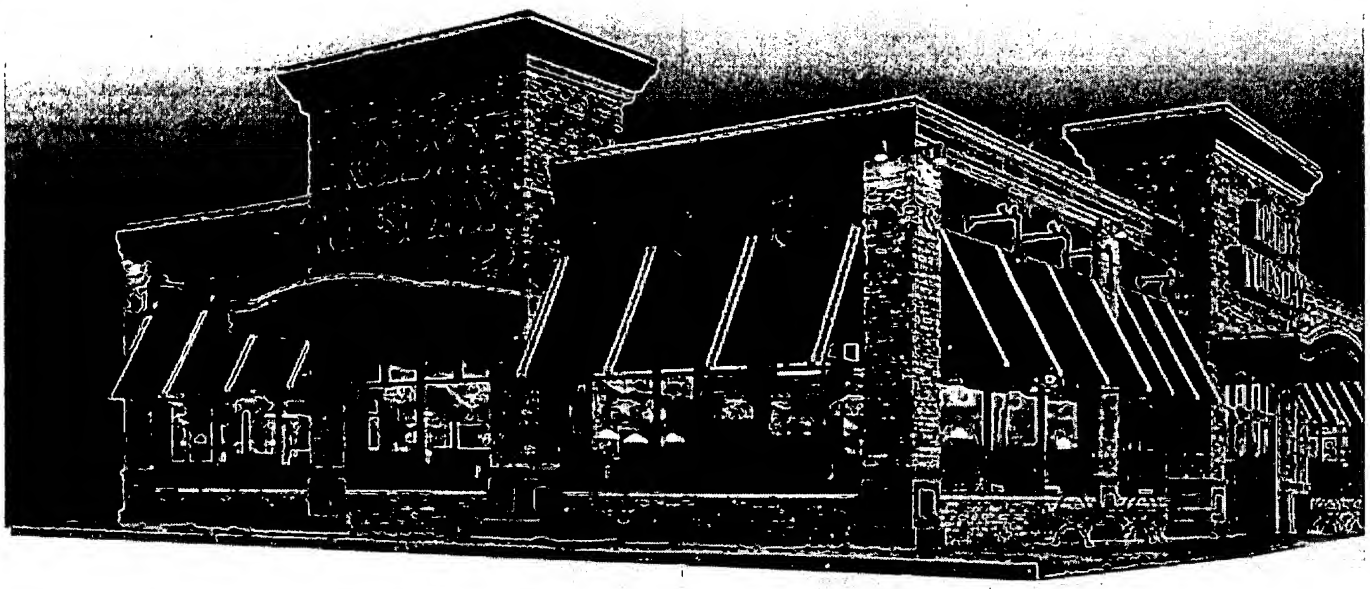

RESTAURANT LOCATOR

Eisenhower Ave

210 SWAMP FOX RD
ALEXANDRIA,
VA 22314
US
Phone: 703-960-1600
Fax: 703-960-5016

To Go Menu

driving directions

address or intersection

city or town

state zip code

Find

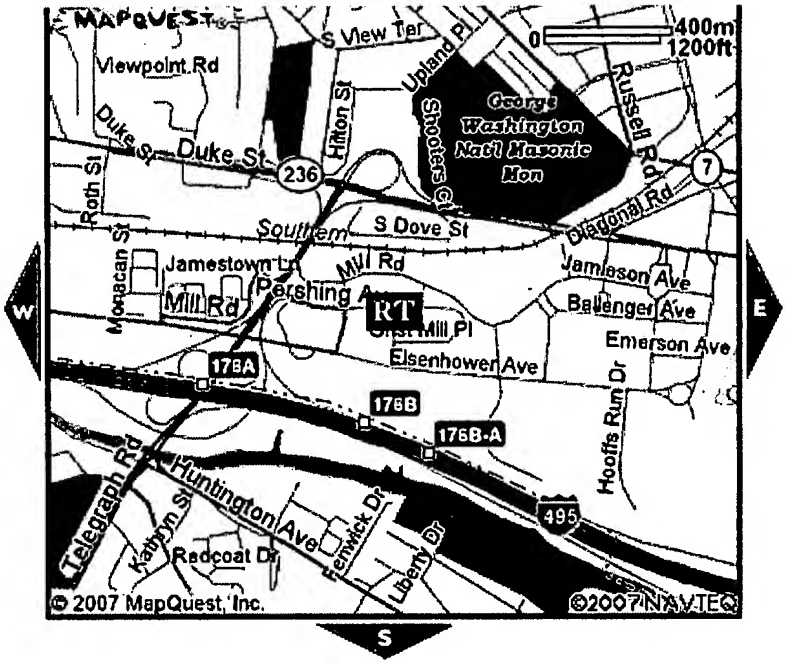

Terms used
memory or **cache** and **partition** or **partitioning** and **shared** and **dedicated** or **exclusive** or **private** and **process**

Sort results by publication date
Display results expanded form
Save results to a Binder
Search Tips
Open results in a new window
Try an Advanced Search
Try this search in The ACM Guic

Results 61 - 80 of 200 Result page: previous 1 2 3 **4** 5 6 7 8 9 10 next
Best 200 shown Relevance scale

61 Energy-aware design of embedded memories: A survey of technologies, architectures, and optimization techniques
Luca Benini, Alberto Macii, Massimo Poncino
February 2003 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 2 Issue 1
Publisher: ACM Press
Full text available: pdf(288.44 KB) Additional Information: full citation, abstract, references, citings, index terms

Embedded systems are often designed under stringent energy consumption budgets, to limit he generation and battery size. Since memory systems consume a significant amount of energy to and to forward data, it is then imperative to balance power consumption and performance in me system design. Contemporary system design focuses on the trade-off between performance and energy consumption in processing and storage units, as well as in their interconnections. Althou memory design is as ...

Keywords: Embedded systems, embedded memories, integration, memories, nonvolatile, syst on-a-chip, volatile

62 Decentralized storage systems: Ivy: a read/write peer-to-peer file system
Athicha Muthitacharoen, Robert Morris, Thomer M. Gil, Benjie Chen
December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI
Publisher: ACM Press
Full text available: pdf(1.65 MB) Additional Information: full citation, abstract, references

Ivy is a multi-user read/write peer-to-peer file system. Ivy has no centralized or dedicated components, and it provides useful integrity properties without requiring users to fully trust eith the underlying peer-to-peer storage system or the other users of the file system.An Ivy file syst consists solely of a set of logs, one log per participant. Ivy stores its logs in the DHash distribut hash table. Each participant finds data by consuiting all logs, but performs modifications by appendi ...

63 Session 1: Nonatomic mutual exclusion with local spinning
James H. Anderson, Yong-Jik Kim
July 2002 **Proceedings of the twenty-first annual symposium on Principles of distributed computing PODC '02**
Publisher: ACM Press

Full text available: pdf(1.12 MB) Additional Information: full citation, abstract, references, citings

We present an *N*-process local-spin mutual exclusion algorithm, based on nonatomic reads and writes, in which each process performs Θ(log *N*) remote memory references to enter and exit its critical section. No atomic read/write algorithm with better asymptotic worst-case time complexity currently known. This suggests that atomic memory is *not* fundamentally required if one is interested in *worst-case* time complexity. The same cannot be said if one is interested ...

64 The state of the art in locally distributed Web-server systems

Valeria Cardellini, Emiliano Casalicchio, Michele Colajanni, Philip S. Yu
June 2002 **ACM Computing Surveys (CSUR)**, Volume 34 Issue 2

Publisher: ACM Press

Full text available: pdf(1.41 MB) Additional Information: full citation, abstract, references, citings, index terms

The overall increase in traffic on the World Wide Web is augmenting user-perceived response time from popular Web sites, especially in conjunction with special events. System platforms that do replicate information content cannot provide the needed scalability to handle large traffic volume and to match rapid and dramatic changes in the number of clients. The need to improve the performance of Web-based services has produced a variety of novel content delivery architecture This article w ...

Keywords: Client/server, World Wide Web, cluster-based architectures, dispatching algorithms distributed systems, load balancing, routing mechanisms

65 On balancing the load in a clustered web farm

Joel L. Wolf, Philip S. Yu
November 2001 **ACM Transactions on Internet Technology (TOIT)**, Volume 1 Issue 2

Publisher: ACM Press

Full text available: pdf(612.40 KB) Additional Information: full citation, abstract, references, citings, index terms

In this article we propose a novel, yet practical, scheme which attempts to optimally balance the on the servers of a clustered Web farm. The goal in solving this performance problem is to achieve minimal average response time for customer requests, and thus ultimately achieve maximal customer throughput. The article decouples the overall problem into two related but distinct mathematical subproblems, one static and one dynamic. We believe this natural decoupling is o the major contrib ...

Keywords: Clustered Web farms, combinatorial optimization, load balancing, resource allocatic problems

66 Computing curricula 2001

September 2001 **Journal on Educational Resources in Computing (JERIC)**

Publisher: ACM Press

Full text available: pdf(613.63 KB) html(2.78 KB) Additional Information: full citation, references, citings, index terms

67 Parallel execution of prolog programs: a survey

Gopal Gupta, Enrico Pontelli, Khayri A.M. Ali, Mats Carlsson, Manuel V. Hermenegildo
July 2001 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 2 4

Publisher: ACM Press

Full text available: pdf(1.95 MB) Additional Information: full citation, abstract, references, citings, index terms

Since the early days of logic programming, researchers in the field realized the potential for exploitation of parallelism present in the execution of logic programs. Their high-level nature, th presence of nondeterminism, and their referential transparency, among other characteristics, m logic programs interesting candidates for obtaining speedups through parallel execution. At the time, the fact that the typical applications of logic programming frequently involve irregular computatio ...

Keywords: Automatic parallelization, constraint programming, logic programming, parallelism, prolog

68 Multiplex: unifying conventional and speculative thread-level parallelism on a chip multiprocessor

Chong-Liang Ooi, Seon Wook Kim, Il Park, Rudolf Eigenmann, Babak Falsafi, T. N. Vijaykumar
June 2001 **Proceedings of the 15th international conference on Supercomputing ICS '01**
Publisher: ACM Press
Full text available: pdf(155.15 KB) Additional Information: full citation, abstract, references, citings, index terms

Recent proposals for Chip Multiprocessors (CMPs) advocate speculative, or implicit, threading in which the hardware employs prediction to peel off instruction sequences (i.e., implicit threads) f the sequential execution stream and speculatively executes them in parallel on multiple process cores. These proposals augment a conventional multiprocessor, which employs explicit threadin with the ability to handle implicit threads. Current proposals focus on only implicitly-threaded cc se ...

69 Data and memory optimization techniques for embedded systems

P. R. Panda, F. Catthoor, N. D. Dutt, K. Danckaert, E. Brockmeyer, C. Kulkarni, A. Vandercappelle, Kjeldsberg
April 2001 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volur Issue 2
Publisher: ACM Press
Full text available: pdf(339.91 KB) Additional Information: full citation, abstract, references, citings, index terms

We present a survey of the state-of-the-art techniques used in performing data and memory-re optimizations in embedded systems. The optimizations are targeted directly or indirectly at the memory subsystem, and impact one or more out of three important cost metrics: area, perform and power dissipation of the resulting implementation. We first examine architecture-independe optimizations in the form of code transoformations. We next cover a broad spectrum of optimiz

Keywords: DRAM, SRAM, address generation, allocation, architecture exploration, code transformation, data cache, data optimization, high-level synthesis, memory architecture customization, memory power dissipation, register file, size estimation, survey

70 Accelerating shared virtual memory via general-purpose network interface support

Angelos Bilas, Dongming Jiang, Jaswinder Pal Singh
February 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 1
Publisher: ACM Press
Full text available: pdf(178.88 KB) Additional Information: full citation, abstract, references, index terms, review

Clusters of symmetric multiprocessors (SMPs) are important platforms for high-performance computing. With the success of hardware cache-coherent distributed shared memory (DSM), a l effort has also been made to support the coherent shared-address-space programming model ir software on clusters. Much research has been done in fast communication on clusters and in protocols for supporting software shared memory across them. However, the performance of software virtual memory (SVM) is sti ...

Keywords: applications, clusters, shared virtual memory, system area networks

71 Cache investment: integrating query optimization and distributed data placement

Donald Kossmann, Michael J. Franklin, Gerhard Drasch, Wig Ag
December 2000 **ACM Transactions on Database Systems (TODS)**, Volume 25 Issue 4
Publisher: ACM Press
Full text available: pdf(210.67 KB) Additional Information: full citation, abstract, references, citings, index terms

Emerging distributed query-processing systems support flexible execution strategies in which ea query can be run using a combination of data shipping and query shipping. As in any distributec environment, these systems can obtain tremendous performance and availability benefits by employing dynamic data caching. When flexible execution and dynamic caching are combined, however, a circular dependency arises: Caching occurs as a by-product of query operator placer but query operator pl ...

Keywords: cache investment, caching, client-server database systems, data shipping, dynamic placement, query optimization, query shipping

72 Parallel shared-memory simulator performance for large ATM networks

Brian Unger, Zhonge Xiao, John Cleary, Jya-Jang Tsai, Carey Williamson
October 2000 **ACM Transactions on Modeling and Computer Simulation (TOMACS)**, Volume 10
Publisher: ACM Press
Full text available: pdf(223.11 KB) Additional Information: full citation, abstract, references, citings, index terms review

A performance comparison between an optimistic and a conservative parallel simulation kernel i presented. Performance of the parallel kernels is also compared to a central-event-list sequentia kernel. A spectrum of ATM network and traffic scenarios representative of those used by ATM networking researchers are used for the comparison. Experiments are conducted with a cell-lev ATM network simulator and an 18-processor SGI PowerChallenge shared-memory multiprocesso The resul ...

Keywords: ATM network modeling, conservative synchronization, optimistic synchronization, p discrete event simulation, time warp

73 Comparative study of page-based and segment-based software DSM through compiler optimization

Junpei Niwa, Takashi Matsumoto, Kei Hiraki
May 2000 **Proceedings of the 14th international conference on Supercomputing ICS '00**
Publisher: ACM Press
Full text available: pdf(1.22 MB) Additional Information: full citation, abstract, references, index terms

The experimental results clearly show that the performance of ADSM scheme is limited by the communication of unnecessary data, while that of the UDSM scheme is limited by the instrumentation overhead. The UDSM scheme reduces transmission of unnecessary data and automatically prevents the severe false sharing at fetch-on-write, which is the problem in the pa based scheme.

74 Multigrain shared memory

Donald Yeung, John Kubiatowicz, Anant Agarwal
May 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 2
Publisher: ACM Press
Additional Information: full citation, abstract, references, citings, index terms

Full text available: pdf(369.18 KB) review

Parallel workstations, each comprising tens of processors based on shared memory, promise co: effective scalable multiprocessing. This article explores the coupling of such small- to medium-s shared-memory multiprocessors through software over a local area network to synthesize large shared-memory systems. We call these systems Distributed Shared-memory MultiProcessors (DSMPs). This article introduces the design of a shared-memory system that uses multiple granularities of sharing, ca ...

Keywords: distributed memory, symmetric multiprocessors, system of systems

75 Piranha: a scalable architecture based on single-chip multiprocessing

Luiz André Barroso, Kourosh Gharachorloo, Robert McNamara, Andreas Nowatzyk, Shaz Qadeer, B: Sano, Scott Smith, Robert Stets, Ben Verghese

May 2000 **ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture ISCA '00**, Volume 28 Issue 2

Publisher: ACM Press

Full text available: pdf(191.10 KB) Additional Information: full citation, abstract, references, citings, index terms

The microprocessor industry is currently struggling with higher development costs and longer d times that arise from exceedingly complex processors that are pushing the limits of instruction-parallelism. Meanwhile, such designs are especially ill suited for important commercial applicati such as on-line transaction processing (OLTP), which suffer from large memory stall times and exhibit little instruction-level parallelism. Given that commercial applications constitute by fa ...

76 System-level power optimization: techniques and tools

Luca Benini, Giovanni de Micheli

April 2000 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volur Issue 2

Publisher: ACM Press

Full text available: pdf(385.22 KB) Additional Information: full citation, abstract, references, citings, index terms

This tutorial surveys design methods for energy-efficient system-level design. We consider elect sytems consisting of a hardware platform and software layers. We consider the three major constituents of hardware that consume energy, namely computation, communication, and stora units, and we review methods of reducing their energy consumption. We also study models for analyzing the energy cost of software, and methods for energy-efficient software design and compilation. This survery ...

77 Atomic heap transactions and fine-grain interrupts

Olin Shivers, James W. Clark, Roland McGrath

September 1999 **ACM SIGPLAN Notices , Proceedings of the fourth ACM SIGPLAN internatio conference on Functional programming ICFP '99**, Volume 34 Issue 9

Publisher: ACM Press

Full text available: pdf(1.45 MB) Additional Information: full citation, abstract, references, citings, index terms

Languages such as Java, ML, Scheme, and Haskell provide automatic storage management, tha garbage collection. The two fundamental operations performed on a garbage-collected heap are "allocate" and "collect." Because the heap is in an inconsistent state during these operations, th must be performed atomically. Otherwise, a heap client might access the heap during a time wh its fundamental invariants do not hold, corrupting the heap.Standard techniques for providing tl atomicity guaran ...

78 Ace: a language for parallel programming with customizable protocols

Mukund Raghavachari, Anne Rogers

August 1999 **ACM Transactions on Computer Systems (TOCS)**, Volume 17 Issue 3

Publisher: ACM Press
Full text available: pdf(297.50 KB) Additional Information: full citation, abstract, references, index terms, review

Customizing the protocols that manage accesses to different data structures within an applicatic improve the performance of software shared-memory programs substantially. Existing systems using customizable protocols are hard to use directly because the mechanisms they provide for manipulating protocols are low-level ones. This article is an in-depth study of the issues involve providing language support for application-specific protocols. We describe the design and implementat ...

Keywords: parallel processing

79 Compile/run-time support for threaded MPI execution on multiprogrammed shared memor machines
Hong Tang, Kai Shen, Tao Yang
May 1999 **ACM SIGPLAN Notices , Proceedings of the seventh ACM SIGPLAN symposium Principles and practice of parallel programming PPoPP '99**, Volume 34 Issue 8
Publisher: ACM Press
Full text available: pdf(1.54 MB) Additional Information: full citation, abstract, references, citings, index terms

MPI is a message-passing standard widely used for developing high-performance parallel applications. Because of the restriction in the MPI computation model, conventional implementa on shared memory machines map each MPI node to an OS process, which suffers serious performance degradation in the presence of multiprogramming, especially when a space/time sl policy is employed in OS job scheduling. In this paper, we study compile-time and run-time sup for MPI by using threads and dem ...

80 Thread scheduling for out-of-core applications with memory server on multicomputers
Yuanyuan Zhou, Limin Wang, Douglas W. Clark, Kai Li
May 1999 **Proceedings of the sixth workshop on I/O in parallel and distributed systems IOPADS '99**
Publisher: ACM Press
Full text available: pdf(861.70 KB) Additional Information: full citation, references, index terms

CiteSeer Find: (memory or cache) and (partition or Documents Citations

Searching for **(memory or cache) and (partition or partitioning) and shared and (dedicated or exclusive or private) and processor**.
Restrict to: Header Title Order by: Expected citations Hubs Usage Date Try: Google (CiteSeer) Google (Web) Yahoo! MSN CSB DBLP
269 documents found. **Order: number of citations.**

Shared Memory Consistency Models: A Tutorial - Adve, Gharachorloo (1995) (Correct) (163 citations)
E M B E R 1 9 9 5 WRL Research Report 95/7 **Shared Memory** Consistency Models: A Tutorial Sarita V. Adve
casaturn.kaist.ac.kr/~sikang/course/CS614/AG95.ps.gz

One or more of the query terms is very common - only partial results have been returned. Try Google (CiteSeer).

The Network Architecture of the Connection Machine CM-5 - Leiserson, Abuhamdeh.. (1994) (Correct) (162 citations)
like many MIMD machines, the CM-5 is a distributed **memory** machine (as opposed to **shared memory** machine [6,
1986. 6] M. Dubois and S. Thakkar, editors. **Cache** Architectures in Tightly Coupled Multiprocessors.
errors. The system operates as one or more user **partitions**. Each **partition** consists of a control
www.ifi.uio.no/~oddvar/bib/papers/Lei+92.ps

The DASH Prototype: Logic Overhead and Performance - Lenoski, Laudon, Joe.. (1993) (Correct) (92 citations)

is that it is feasible to build large-scale **shared-memory** multiprocessors with hardware **cache** coherence.
www-flash.stanford.edu/pub/flash/TPDS93.ps.Z

Speculative Versioning Cache - Gopal (1998) (Correct) (67 citations)
execution of a sequential program. Such ambiguous **memory** dependences can be overcome by **memory** dependence
Speculative Versioning **Cache** Sridhar Gopal y T.N.Vijaykumar James E.
ftp.cs.wisc.edu/sohi/papers/1998/hpca.svc.ps.gz

Adaptive Parallelism and Piranha - Carriero, Freeman, Gelernter.. (1995) (Correct) (60 citations)
space is a virtual **shared**, associative, object **memory** accessible to all nodes within a parallel
user jobs are statically assigned to some **partition**. But there are problems with this approach. If
demands that tend to be imposed at most sites on **shared** resources. Further, cheap hardware, economies of
ftp.cs.yale.edu/WWW/HTML/YALE/CS/Linda/papers/shortp.ps

Techniques for Reducing Consistency-Related.. - Carter, Bennett.. (1993) (Correct) (59 citations)
Communication in Distributed **Shared Memory** Systems John B. Carter, John K. Bennett and
mancos.cs.utah.edu/papers/munin.ps.Z

Synthesis and Simulation of Digital Systems.. - Gupta, Coelho, Jr.. (1992) (Correct) (55 citations)
Tasks, And It Is Not Yet An Automated Tool. Asic **Memory** MI Program User Data Interface Buffer
description as input. The input system model is **partitioned** into hardware and software components based
description as input. The input system model is **partitioned** into hardware and software components based on
akebono.stanford.edu/users/coelho/papers/dac-92.ps.gz

Converting Thread-Level Parallelism to.. - Lo, Eggers, Emer.. (1997) (Correct) (43 citations)
parallelism or multithreading: interference in the **memory** system and branch prediction hardware. We find
threads cause interthread interference in the **caches** and place greater demands on the **memory** system,
www.cs.washington.edu/research/smt/papers/tlp2ilp.final.ps

Thread Scheduling for Multiprogrammed Multiprocessors - Arora, Blumofe, Plaxton (1998) (Correct) (41 citations)
present a user-level thread scheduler for **shared-memory** multiprocessors, and we analyze its performance
in particular those that employ static space **partitioning** [13, 26] or coscheduling [15, 26, 29]they

in particular those that employ static space **partition**ing [13, 26] or coscheduling [15, 26, 29]they do
www.cs.utexas.edu/users/rdb/papers/SPAA98.ps.gz

Dynamic Coscheduling on Workstation Clusters - Patrick Sobalvarro (1998) (Correct) (40 citations)
access to input/output devices ffl coordinated **memory** management ffl efficient parallel computing with coscheduling and process control (dynamic space-**partition**ing) performed similarly in the experiments coscheduling and process control (dynamic space-**partition**ing) performed similarly in the experiments
www-csag.ucsd.edu/papers/csag/external/fmdcs-sosp97.ps

Robust Partitioning Policies of Multiprocessor Systems - Rosti, Smirni, Dowdy.. (1993) (Correct) (40 citations)
to the underlying hardware platform. In **shared memory** environments, dynamic schemes may be possible
Robust **Partition**ing Policies of Multiprocessor Systems E.
Robust **Partition**ing Policies of Multiprocessor Systems E. Rosti
www.cs.wm.edu/~esmirni/docs/perfeval.ps.gz

Scheduling Large-Scale Parallel Computations on Networks of.. - Blumofe (1994) (Correct) (34 citations)
following goals: 1. Preserve communications and **memory** locality. 2. Accommodate dynamic parallelism.
the overhead incurred by network communication, **cache** misses, and page faults becomes ever more
its **processors** into a small number of fixedsize **partitions**. Each **partition** is run either in **dedicated**
www.cs.utexas.edu/users/rdb/papers/HPDC94.ps.gz

Paging Tradeoffs in Distributed-Shared-Memory.. - Burger, Hyder, Miller, Wood (1994) (Correct) (33 citations)
Paging Tradeoffs In Distributed-**Shared-Memory** Multiprocessors Douglas C. Burger And Rahmat
ftp.cs.wisc.edu/wwt/jsc96_paging.ps

Coordinated Checkpointing-Rollback Error Recovery for.. - Janakiraman, Tamir (1994) (Correct) (32 citations)
Error Recovery for Distributed **Shared Memory** Multicomputers G. Janakiraman and Yuval Tamir
where **memory** is physically distributed, using **cache** coherency protocols [3, 11, 13]The reliability
tasks simultaneously. Unique task identifiers **partition** the total virtual system space into disjoint
ftp.cs.ucla.edu/tech-report/94-reports/940027.ps.Z

Efficient Support for Multicomputing on ATM Networks - Thekkath, Levy, Lazowska (1993) (Correct) (31 citations)
communication model based on the notion of remote **memory** access. Applications executing on one host can
I/O bus, using network interfaces connected to the **cache** bus instead, as is done in **dedicated**
have favored simpler models that involve **partition**ing data among **processors** and more direct
ftp.cs.washington.edu/tr/1993/04/UW-CSE-93-04-03.PS.Z

Parallel Job Scheduling: Issues and Approaches - Feitelson, Rudolph (1995) (Correct) (26 citations)
of much larger aggregate **cache** capacity, physical **memory** size, and I/O bandwidth. Of course, when an
shorter time, and make use of much larger aggregate **cache** capacity, physical **memory** size, and I/O
are to use a global queue, use variable **partition**ing, use dynamic **partition**ing, and use gang
www.cs.huji.ac.il/~feit/parsched/p-95-1.ps.gz

First 20 documents Next 20

Try your query at: Google (CiteSeer) Google (Web) Yahoo! MSN CSB DBLP

Results for "(((memory or cache) and (partition or partitioning) and (shared) and (dedicated or exclusive or priv..." e-mail

Your search matched **12** of **1540244** documents.

A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

1. **Customized on-chip memories for embedded chip multiprocessors**
Ozturk, O.; Kandemir, M.; Chen, G.; Irwin, M.J.; Karakoy, M.;
Design Automation Conference, 2005. Proceedings of the ASP-DAC 2005. Asi Pacific
Volume 2, 18-21 Jan. 2005 Page(s):743 - 748 Vol. 2
Digital Object Identifier 10.1109/ASPDAC.2005.1466448
AbstractPlus | Full Text: PDF(381 KB) IEEE CNF
Rights and Permissions

2. **A hierarchical processor scheduling policy for distributed-memory multi systems**
Dandamudi, S.P.; Thyagaraj, T.K.;
High Performance Computing, 1997. Proceedings. Fourth International Confer
18-21 Dec. 1997 Page(s):218 - 223
Digital Object Identifier 10.1109/HIPC.1997.634496
AbstractPlus | Full Text: PDF(756 KB) IEEE CNF
Rights and Permissions

3. **Performance sensitivity of space-sharing processor scheduling in distrib multicomputers**
Dandamudi, S.P.; Yu, H.;
Parallel Processing Symposium, 1998. 1998 IPPS/SPDP. Proceedings of the I International...and Symposium on Parallel and Distributed Processing 1998
30 March-3 April 1998 Page(s):403 - 409
Digital Object Identifier 10.1109/IPPS.1998.669948
AbstractPlus | Full Text: PDF(728 KB) IEEE CNF
Rights and Permissions

4. **Multi-level on-chip memory hierarchy design for embedded chip multipro**
Ozturk, O.; Kandemir, M.; Irwin, M.J.; Tosun, S.;
Parallel and Distributed Systems, 2006. ICPADS 2006. 12th International Conf
Volume 1, 12-15 July 2006 Page(s):8 pp.
Digital Object Identifier 10.1109/ICPADS.2006.66
AbstractPlus | Full Text: PDF(368 KB) IEEE CNF
Rights and Permissions

5. **Compositional, efficient caches for a chip multi-processor**

» Key

IEEE JNL	IEEE Journal or Magazine
IET JNL	IET Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IET CNF	IET Conference Proceeding
IEEE STD	IEEE Standard

Molnos, A.M.; Heijligers, M.J.M.; Cotofana, S.D.; van Eijndhoven, J.T.J.;
Design, Automation and Test in Europe, 2006. DATE '06. Proceedings
Volume 1, 6-10 March 2006 Page(s):1 - 6

AbstractPlus | Full Text: PDF(296 KB) IEEE CNF
Rights and Permissions

6. **Iteration Partitioning for Resolving Stride Conflicts on Cache-Coherent M**
Karen A. Tomko; Santosh G. Abraham;
Parallel Processing, 1993. ICPP 1993 Volume 2. International Conference on
Volume 2, Aug. 1993 Page(s):95 - 102
Digital Object Identifier 10.1109/ICPP.1993.107

AbstractPlus | Full Text: PDF(1499 KB) IEEE CNF
Rights and Permissions

7. **A combined method for maintaining large indices in multiprocessor multi environments**
Matsliach, G.; Shmueli, O.;
Knowledge and Data Engineering, IEEE Transactions on
Volume 6, Issue 3, June 1994 Page(s):479 - 496
Digital Object Identifier 10.1109/69.334867

AbstractPlus | Full Text: PDF(1608 KB) IEEE JNL
Rights and Permissions

8. **OpenMP Directive Extension for BlackFin 561 Dual Core Processor**
Hee Seo; Seon Wook Kim;
Computer and Information Technology, 2006. CIT '06. The Sixth IEEE Internat on
Sept. 2006 Page(s):49 - 49
Digital Object Identifier 10.1109/CIT.2006.131

AbstractPlus | Full Text: PDF(720 KB) IEEE CNF
Rights and Permissions

9. **Implementing malleability on MPI jobs**
Utrera, G.; Corbalan, J.; Labarta, J.;
Parallel Architecture and Compilation Techniques, 2004. PACT 2004. Proceed International Conference on
29 Sept.-3 Oct. 2004 Page(s):215 - 224
Digital Object Identifier 10.1109/PACT.2004.1342555

AbstractPlus | Full Text: PDF(289 KB) IEEE CNF
Rights and Permissions

10. **ETA: experience with an Intel Xeon processor as a packet processing en**
Regnier, G.; Minturn, D.; McAlpine, G.; Saletore, V.A.; Foong, A.;
Micro, IEEE
Volume 24, Issue 1, Jan.-Feb. 2004 Page(s):24 - 31
Digital Object Identifier 10.1109/MM.2004.1268989

AbstractPlus | References | Full Text: PDF(267 KB) IEEE JNL
Rights and Permissions

11. **Platform-based design from parallel C specifications**
Auge, I.; Petrot, F.; Donnet, F.; Gomez, P.;
Computer-Aided Design of Integrated Circuits and Systems, IEEE Transaction
Volume 24, Issue 12, Dec. 2005 Page(s):1811 - 1826
Digital Object Identifier 10.1109/TCAD.2005.852431

AbstractPlus | Full Text: PDF(816 KB) IEEE JNL
Rights and Permissions

12. **Implementing MPI under AP/Linux**
Sitsky, D.; Mackerras, P.; Tridgell, A.; Walsh, D.;

MPI Developer's Conference, 1996. Proceedings., Second
1-2 July 1996 Page(s):32 - 39
Digital Object Identifier 10.1109/MPIDC.1996.534092

Results for "(((memory or cache) and (dynamic or dynamically) and (partition or partitioning or allocate or alloc..." e-mail

Your search matched **111** of **1540244** documents.

A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

1. **Data caches in multitasking hard real-time systems**
   Vera, X.; Lisper, B.; Jingling Xue;
   Real-Time Systems Symposium, 2003. RTSS 2003. 24th IEEE
   2003 Page(s):154 - 165
   Digital Object Identifier 10.1109/REAL.2003.1253263
   AbstractPlus | Full Text: PDF(364 KB) IEEE CNF
   Rights and Permissions

2. **A configurable logic architecture for dynamic hardware/software partitior**
   Lysecky, R.; Vahid, F.;
   Design, Automation and Test in Europe Conference and Exhibition, 2004. Proc
   Volume 1, 16-20 Feb. 2004 Page(s):480 - 485 Vol.1
   Digital Object Identifier 10.1109/DATE.2004.1268892
   AbstractPlus | Full Text: PDF(243 KB) IEEE CNF
   Rights and Permissions

3. **Adaptive cache compression for high-performance processors**
   Alameldeen, A.R.; Wood, D.A.;
   Computer Architecture, 2004. Proceedings. 31st Annual International Symposi
   19-23 June 2004 Page(s):212 - 223
   Digital Object Identifier 10.1109/ISCA.2004.1310776
   AbstractPlus | Full Text: PDF(324 KB) IEEE CNF
   Rights and Permissions

4. **Affinity scheduling of unbalanced workloads**
   Subramaniam, S.; Eager, D.L.;
   Supercomputing '94. Proceedings
   14-18 Nov. 1994 Page(s):214 - 226
   Digital Object Identifier 10.1109/SUPERC.1994.344281
   AbstractPlus | Full Text: PDF(836 KB) IEEE CNF
   Rights and Permissions

5. **A memory-layout oriented run-time technique for locality optimization on**
   Yong Yan; Xiaodong Zhang; Zhao Zhang;
   Parallel Processing, 1998. Proceedings. 1998 International Conference on
   10-14 Aug. 1998 Page(s):189 - 196
   Digital Object Identifier 10.1109/ICPP.1998.708484

AbstractPlus | Full Text: PDF(96 KB) IEEE CNF
Rights and Permissions

6. **On dynamic speculative thread partitioning and the MEM-slicing algorithı**
Codrescu, L.; Wills, D.S.;
Parallel Architectures and Compilation Techniques, 1999. Proceedings. 1999 I
Conference on
12-16 Oct. 1999 Page(s):40 - 46
Digital Object Identifier 10.1109/PACT.1999.807404
AbstractPlus | Full Text: PDF(112 KB) IEEE CNF
Rights and Permissions

7. **Hardware support for real-time embedded multiprocessor system-on-a-cl management**
Shalan, M.; Mooney, V.J., III;
Hardware/Software Codesign, 2002. CODES 2002. Proceedings of the Tenth I
Symposium on
6-8 May 2002 Page(s):79 - 84
Digital Object Identifier 10.1109/CODES.2002.1003605
AbstractPlus | Full Text: PDF(438 KB) IEEE CNF
Rights and Permissions

8. **Dynamic partitioning of processing and memory resources in embedded architectures**
Xue, L.; Ozturk, O.; Li, F.; Kandemir, M.; Kolcu, I.;
Design, Automation and Test in Europe, 2006. DATE '06. Proceedings
Volume 1, 6-10 March 2006 Page(s):6 pp.
AbstractPlus | Full Text: PDF(184 KB) IEEE CNF
Rights and Permissions

9. **Energy-efficient physically tagged caches for embedded processors with**
Petrov, P.; Tracy, D.; Orailoglu, A.;
Design Automation Conference, 2005. Proceedings. 42nd
13-17 June 2005 Page(s):17 - 22
AbstractPlus | Full Text: PDF(456 KB) IEEE CNF
Rights and Permissions

10. **Fair cache sharing and partitioning in a chip multiprocessor architecture**
Kim, S.; Chandra, D.; Solihin, Y.;
Parallel Architecture and Compilation Techniques, 2004. PACT 2004. Proceed
International Conference on
29 Sept.-3 Oct. 2004 Page(s):111 - 122
Digital Object Identifier 10.1109/PACT.2004.1342546
AbstractPlus | Full Text: PDF(442 KB) IEEE CNF
Rights and Permissions

11. **Dynamic scheduling strategies for shared-memory multiprocessors**
Hamidzadeh, B.; Lilja, D.J.;
Distributed Computing Systems, 1996., Proceedings of the 16th International (
27-30 May 1996 Page(s):208 - 215
Digital Object Identifier 10.1109/ICDCS.1996.507918
AbstractPlus | Full Text: PDF(764 KB) IEEE CNF
Rights and Permissions

12. **Energy-efficient cache architecture for multimedia applications**
Chia-Lin Yang; Chien-hao Lee; Hung-Wei Tseng;
Emerging Information Technology Conference, 2005.
15-16 Aug. 2005 Page(s):2 pp.
Digital Object Identifier 10.1109/EITC.2005.1544355

AbstractPlus | Full Text: PDF(104 KB) IEEE CNF
Rights and Permissions

☐ **13. A Reconfigurable Cache Architecture for Object-Oriented Embedded Sys**
Modarressi, M.; Hessabi, S.; Goudarzi, M.;
Electrical and Computer Engineering, Canadian Conference on
May 2006 Page(s):959 - 962
Digital Object Identifier 10.1109/CCECE.2006.277819
AbstractPlus | Full Text: PDF(95 KB) IEEE CNF
Rights and Permissions

☐ **14. Performance sensitivity of space-sharing processor scheduling in distrit multicomputers**
Dandamudi, S.P.; Yu, H.;
Parallel Processing Symposium, 1998. 1998 IPPS/SPDP. Proceedings of the I International...and Symposium on Parallel and Distributed Processing 1998
30 March-3 April 1998 Page(s):403 - 409
Digital Object Identifier 10.1109/IPPS.1998.669948
AbstractPlus | Full Text: PDF(728 KB) IEEE CNF
Rights and Permissions

☐ **15. Memory management techniques for time warp on a distributed memory**
Preiss, B.R.; Loucks, W.M.;
Parallel and Distributed Simulation, 1995. (PADS'95), Proceedings., Ninth Wo No.95TB8096)
14-16 June 1995 Page(s):30 - 39
Digital Object Identifier 10.1109/PADS.1995.404319
AbstractPlus | Full Text: PDF(724 KB) IEEE CNF
Rights and Permissions

☐ **16. Layout-driven memory synthesis for embedded systems-on-chip**
Benini, L.; Macchiarulo, L.; Macii, A.; Poncino, M.;
Very Large Scale Integration (VLSI) Systems, IEEE Transactions on
Volume 10, Issue 2, April 2002 Page(s):96 - 105
Digital Object Identifier 10.1109/92.994985
AbstractPlus | References | Full Text: PDF(301 KB) IEEE JNL
Rights and Permissions

☐ **17. Custom design of multi-level dynamic memory management subsystem f systems**
Mamagkakis, S.; Atienza, D.; Poucet, C.; Catthoor, F.; Soudris, D.; Mendias, J
Signal Processing Systems, 2004. SIPS 2004. IEEE Workshop on
2004 Page(s):170 - 175
Digital Object Identifier 10.1109/SIPS.2004.1363044
AbstractPlus | Full Text: PDF(411 KB) IEEE CNF
Rights and Permissions

☐ **18. Hardware-based methods for maintaining cache coherency in a multipro**
Jones, J.; Haggard, R.L.;
System Theory, 1998. Proceedings of the Thirtieth Southeastern Symposium c
8-10 March 1998 Page(s):549 - 553
Digital Object Identifier 10.1109/SSST.1998.660134
AbstractPlus | Full Text: PDF(524 KB) IEEE CNF
Rights and Permissions

☐ **19. Memory-centric architectures: why and perhaps what**
Burger, D.; Goodman, J.R.;
Innovative Architecture for Future Generation High-Performance Processors a
22-24 Oct. 1997 Page(s):92

Digital Object Identifier 10.1109/IWIA.1997.670413

AbstractPlus | Full Text: PDF(100 KB) IEEE CNF
Rights and Permissions

☐ **20. Task Partitioning Oriented Architecture Exploration Method for Dynamic Architectures**
Taniguchi, I.; Ueda, K.; Sakanushi, K.; Takeuchi, Y.; Imai, M.;
Very Large Scale Integration, 2006 IFIP International Conference on
Oct. 2006 Page(s):290 - 295
Digital Object Identifier 10.1109/VLSISOC.2006.313249

AbstractPlus | Full Text: PDF(495 KB) IEEE CNF
Rights and Permissions

☐ **21. Microarchitectural wire management for performance and power in partit architectures**
Balasubramonian, R.; Muralimanohar, N.; Ramani, K.; Venkatachalapathy, V.;
High-Performance Computer Architecture, 2005. HPCA-11. 11th International
12-16 Feb. 2005 Page(s):28 - 39
Digital Object Identifier 10.1109/HPCA.2005.21

AbstractPlus | Full Text: PDF(208 KB) IEEE CNF
Rights and Permissions

☐ **22. Reducing transitions on memory buses using sector-based encoding tec**
Aghaghiri, Y.; Fallah, F.; Pedram, M.;
Low Power Electronics and Design, 2002. ISLPED '02. Proceedings of the 20(Symposium on
2002 Page(s):190 - 195

AbstractPlus | Full Text: PDF(773 KB) IEEE CNF
Rights and Permissions

☐ **23. Software-controlled cache architecture for energy efficiency**
Chia-Lin Yang; Hung-Wei Tseng; Chia-Chiang Ho; Ja-Ling Wu;
Circuits and Systems for Video Technology, IEEE Transactions on
Volume 15, Issue 5, May 2005 Page(s):634 - 644
Digital Object Identifier 10.1109/TCSVT.2005.846444

AbstractPlus | References | Full Text: PDF(1040 KB) IEEE JNL
Rights and Permissions

☐ **24. From architecture to layout: partitioned memory synthesis for embedded chip**
Benini, L.; Macchiarulo, L.; Macii, A.; Poncino, M.;
Design Automation Conference, 2001. Proceedings
2001 Page(s):784 - 789

AbstractPlus | Full Text: PDF(636 KB) IEEE CNF
Rights and Permissions

☐ **25. Performance analysis and prediction of processor scheduling strategies multiprogrammed shared-memory multiprocessors**
Yue, K.K.; Lilja, D.J.;
Parallel Processing, 1996., Proceedings of the 1996 International Conference
Volume 3, 12-16 Aug. 1996 Page(s):70 - 78 vol.3
Digital Object Identifier 10.1109/ICPP.1996.538561

AbstractPlus | Full Text: PDF(804 KB) IEEE CNF
Rights and Permissions

View: **1-25** | 26-5